arXiv:1402.7118v1 [cs.CR] 28 Feb 2014

Lightweight Self-Bootstrapping Multiparty Computations of Time-Series Data with Custom Collusion Tolerance

Michael Clear*, Constantinos Patsakis, and Paul Laird

School of Computer Science and Statistics,
Trinity College Dublin

Abstract. In this work we compare two recent multiparty computation (MPC) protocols for private summation in terms of performance. Both protocols allow multiple rounds of aggregation from the same set of public keys generated by parties in an initial stage. We instantiate the protocols with a fast elliptic curve and provide an experimental comparison of their performance for different phases of the protocol. Furthermore, we introduce a technique that allows the computational load of both protocols to be reduced at the expense of protection against collusion tolerance. We prove that both protocols remain secure with this technique, and evaluate its impact on collusion tolerance and the number of rounds supported.

Keywords: multiparty computation, private summation, custom collusion tolerance, private aggregation

1 Introduction

Modern computing has reached a point that allows users from every part of the globe to exchange information seamlessly. Nevertheless, the environment cannot be considered friendly. Cyber attacks each year are reaching a new peak, while recently disclosed events clearly indicate that privacy measures are not properly deployed or applied. The hostile environment to which modern users are exposed to has triggered the generation of many security and privacy protocols. In many scenarios, users have to co-operate to perform several tasks which result in some of their data being partially or fully exposed. Since the latter might not be an acceptable option, given the sensitivity of the submitted information, there has been much research effort in the area of secure multi-party computation (MPC).

While there has been much progress made in recent years towards making MPC practical, the efficiency of many protocols is still not acceptable for many real-world applications. To cover this gap, apart from general MPC protocols (which allow any function to be computed), many MPC protocols for specialized computations have been introduced. While they might support only a small

[1]The author's work is funded by the Irish Research Council EMBARK Initiative.

set of functions, these protocols can be obtain much improved performance, even in devices with limited computing resources. Thus, apart from the interest in general MPC protocols, many application needs are pushing towards the development of more targeted MPC protocols. A typical example of this trend are protocols which allow privacy-aware summation of values. These protocols are widely used for load monitoring of smart meters, privacy-aware participatory sensing, and generally for submitting time series data with privacy.

Recently, a new protocol was proposed by Patsakis, Clear and Laird [26] (referred to here as PCL) that allows privacy-aware multiparty aggregation of values. A notable feature is that it supports multiple rounds of aggregation from the same set of public keys published by the users, and it relies only on the standard Decisional Diffie-Hellman (DDH) assumption. PCL is based on a protocol due to Kursawe, Danezis and Kohlweiss [20] (referred to here as KDK) which also has a variant that supports multiple rounds of aggregation by exploiting bilinear pairings. Both these papers do not consider the performance of their protocols in practice. In fact, [26] does not address its performance relative to multi-round KDK. In this paper, we extensively compare the performance of both protocols, and show that in practice PCL outperforms multi-round KDK by a significant margin, largely due to the cost of pairings, but also due to the larger finite field in which recovery of the sum takes place. Following on from concrete performance results, we propose a new extension that is applicable to both protocols, which leads to considerable performance improvements. In PCL, protection against collusion tolerance is traded off against the number of rounds supported. Our extension trades collusion tolerance further (in both protocols) to achieve a performance gain. However, since a tolerance of $1/3-1$ of the parties is satisfactory for many real-world applications (note that this corresponds to the Byzantine optimum), we argue that reducing tolerance to (say) $1/3$ is justified given the significant gains in performance.

1.1 Main contributions

The main technical contribution of this work are as follows. Firstly, a thorough performance analysis is given of the multi-round protocols PCL and KDK. To the best of our knowledge, this is the first assessment of the practicality of multi-round KDK. Secondly, we introduce a new technique that is applicable to both protocols that allows all parties to reduce their computational load by adjusting their collusion tolerance. We prove security for the extended versions of both protocols that employ this technique.

1.2 Organization of this work

The rest of this work is organized as follows. In the next section we provide a brief overview of the related work in this area, along with a description of the two protocols, PCL and KDK in their single-round and multi round versions. Afterwards, in section 3, we compare the performance of the two protocols.

Sections 4 and five describe methods to reduce the computational load of PCL and KDK protocols respectively.

2 Related work

In many application scenarios users have to co-operate to achieve specific tasks, however, they need to know that the that they will send will not be disclosed. Based on this problem, Yao introduced the concept of secure multiparty computation was [30]. Despite the huge initial advances [31,15,8,4], only recently did real-world implementations become practical [24,3,5].

Due to efficiency, secure multiparty computation can be categorized into schemes that allow arbitrary computations to be performed without leaking information, and more specialized protocols which allow the private evaluation of particular functions, such as summation. In the first category we have schemes based on Garbled Circuits (efficient implementations include [22,23,21]), Nielsen's protocol using Oblivious Transfer [25] and Oblivious RAM [14,17]. Other protocols use arithmetic circuits such as those based on the BGW protocol e.g: VIFF [10] and SPDZ which employs fully homomorphic encryption [11]. As already discussed, the second category includes more application-specific protocols. For instance for calculating the private sum of n users we have the scheme of Clifton *et al.* [9], the two round scheme of Yang *et al.* [29] and the single round scheme of Shi *et al.* but relies on the previously distribution of shares of shares of zeros from a trusted third party [28]. Other protocols are focused on calculating one bit multi-party computations, such as DC-net [7,16] and [18].

3 Preliminaries

3.1 Security Definition

We adopt the standard simulation-based definition of security in the semi-honest model. We base our definition below on Definition 2.1 in [1]. Here we consider only computational security, and relax the more standard definition to deterministic functionalities with a single output. Note that this definition is general enough to accommodate multi-round aggregation.

Let $\boldsymbol{m} \in (\{0,1\}^*)^n$ be a vector of the inputs from each party and let π be a protocol. We define the view of a party P_i in the execution of protocol π with input vector $\boldsymbol{m}$ as

$$\mathsf{VIEW}_i^{\pi}(\boldsymbol{x}) = (m_i, r_i, \mu_i^{(1)}, \ldots, \mu_i^{(\ell)})$$

where m_i is party $P_i's$ input, r_i is its random coins and $\mu_i^{(1)}, \ldots, \mu_i^{(\ell)}$ are the ℓ protocol messages it received during the protocol execution. Similarly, the combined view of a set of $I \subseteq \{1, \ldots, n\}$ parties is denoted by $\mathsf{VIEW}_I^{\pi}(\boldsymbol{x})$.

Definition 1 (t-privacy of n-party protocols for deterministic aggregation functionalities). *Let $f : (\{0,1\}^*)^n \to (\{0,1\}^*)$ be a deterministic n-ary functionality and let π be a protocol that* correctly *computes f. We say that π if t-private if for every $\boldsymbol{m} \in (\{0,1\}^*)^n$ where $|m_1| = \ldots = |m_n|$, there exists a PPT algorithm $\mathcal{S}$ such that for every $I \subset \{1,\ldots,n\}$ with $|I| \leq t$, and every $\boldsymbol{m} \in (\{0,1\})^n$ where $|m_1| = \ldots = |m_n|$, it holds that:*

$$\{\mathsf{VIEW}_I^{\pi}(\boldsymbol{m})\} \underset{C}{\approx} \{\mathcal{S}(I, \boldsymbol{m}_I, f_I(\boldsymbol{m}))\}. \tag{1}$$

where $\underset{C}{\approx}$ denotes computational indistinguishability.

3.2 KDK Single-Round Protocol

Kursawe, Danezis and Kohlweiss (KDK) [20] present a specialized multiparty computation (MPC) protocol for private summation, which is shown is be secure in the semihonest model under the Decisional Diffie-Hellman (DDH) assumption. We refer to this protocol as KDK. In their protocol, n parties $P_1, \ldots, P_n$ can compute a joint sum of their inputs $m_1, \ldots, m_n \in \{0, \ldots, \beta\}$ for some positive integer β. An overview of their protocol follows.

Let p be a prime. The "public parameters" used in the protocol consist of a description of a cyclic group $\mathbb{G}$ of order p together with a generator g of $\mathbb{G}$. It is assumed that DDH is intractable in $\mathbb{G}$. These public parameters $\mathsf{PP} = (\mathbb{G}, g, p)$ are known to all parties P_i. The group operation of $\mathbb{G}$ is written multiplicatively.

1. **Setup:** Party P_i generates a secret key $x_i \in \mathbb{Z}_p$ and computes her public key $u_i = g^{x_i} \in \mathbb{G}$. She broadcasts u_i.
2. For every $r \in \{1, \ldots, \ell\}$:
   Main Round:
   - Party P_i chooses her input $m_i \in \{0, \ldots, \beta\}$.
   - Compute $w \leftarrow \prod_{j\in 1}^{i-1} u_j^{-1} \cdot \prod_{j\in i+1}^{n} u_j \in G$.
   - Compute $v_i \leftarrow w^{x_i} \cdot g^{m_i} \in G$.
   - Broadcast v_i.
3. **Output:** The protocol produces an output in $\{0, \ldots, n\beta\}$, namely the sum of the user inputs. To compute the sum σ:
   - Compute $z \leftarrow \prod_{j=1}^{n} v_j$.
   - Use Pollard's Lambda algorithm to compute the discrete $\log \sigma \in \{0, \ldots, n\beta\}$ of z with respect to g in G. The time complexity of Pollard's lambda algorithm is $\sqrt{n\beta}$.
   - Output σ.

It can be easily observed that

$$\prod_{j=1}^{n} v_j = g^{\sum_{j=1}^{n} m_j} \tag{2}$$

3.3 KDK Multi-Round Protocol

If the protocol must be run a number of times, it would be desirable to avoid re-running the "Setup" phase above which involves each party generating and broadcasting a new public key; in practice, a verification step for these keys may also be needed. To re-use the published keys $u_i, \ldots, u_n$ for more than a single round of aggregation, Kursawe et al. propose an extension of their protocol that facilitates multiple-rounds. In fact, their multi-round protocol accommodates an unbounded number of rounds. They make use of bilinear pairings to achieve this. More details on bilinear pairings are provided in Section 4 when we address practical issues. We give a very brief overview here that is sufficient to understand the multi-round protocol. The following is based on the definition from [12] (Section 2).

Definition 2. *Let $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ be cyclic groups of prime order q. We write $\mathbb{G}_1$ and $\mathbb{G}_2$ additively, and $\mathbb{G}_T$ multiplicatively. A bilinear pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is an efficiently computable map satisfying*

- Bilinearity: *$e(aP, bQ) = e(P, Q)^{ab}$ for all $P \in \mathbb{G}_1$, $Q \in \mathbb{G}_2$ and $a, b \in \mathbb{Z}_q^*$.*
- Non-degeneracy: *If P is a generator for $\mathbb{G}_1$ and Q is a generator for $\mathbb{G}_2$, then $e(P, Q) \neq 1$.*

Examples of bilinear pairings (or their modifications) that are used in cryptography include the Tate Pairing [2,13], Weil Pairing [6], and Ate Pairing [19].

Let $\mathbb{G}_1$, $\mathbb{G}_2$ and $\mathbb{G}_T$ be cyclic groups of prime order p. Let $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ be a cryptographic bilinear pairing meeting the conditions of Definition 2. Furthermore, the Bilinear Decisional Diffie Hellman (BDDH) assumption is expected to hold with respect to $\mathbb{G}_1$, $\mathbb{G}_2$, $\mathbb{G}_T$ and e. Let $H_2 : \mathbb{Z} \to \mathbb{G}_2$ be a hash function. The main changes to KDK to support multiple rounds are as follows (optimizations are discussed later):

- The public parameters include generators $P \in \mathbb{G}_1$, $Q \in \mathbb{G}_2$ and $g = e(P, Q) \in \mathbb{G}_T$.
- The public keys are generated as $U_i \leftarrow x_i P \in \mathbb{G}_1$ for all $1 \leq i \leq n$.
- In round k, party P_i computes
  - $Q_k \leftarrow H_2(k) \in \mathbb{G}_2$ (i.e. for a good choice of H_2, we have $Q_k = rQ$ for some uniformly random r, which is intractable to find).
  - $w \leftarrow \prod_{j \in 1}^{i-1} e(U_j, Q_k)^{-1} \cdot \prod_{j \in i+1}^{n} e(U_j, Q_k) \in \mathbb{G}_T$.

The rest of the protocol remains unchanged except that the computations are performed in $\mathbb{G}_T$, and P_i may choose a different input value in every round. Naturally, the output of the protocol is then $(\sigma_1, \ldots, \sigma_\ell) \in \{0, \ldots, n\beta\}^\ell$ if ℓ rounds are executed.

3.4 PCL Multi-Round Protocol

Patsakis, Clear and Laird [26] introduced another multi-round variant of KDK without pairings. Their protocol (referred to here as PCL) allows a bounded

number of rounds ℓ to be performed from the same public key information. However, ℓ depends on the acceptable collusion tolerance $t \leq n$. Both single-round and multi-round KDK are t-private for any $t \leq n$. On the other hand, in order for PCL to be t-private, at most $\ell = \lfloor \frac{n-t}{2} \rfloor$ rounds are permitted. Concretely, for $t = n/3$ (Byzantine tolerance) and $n = 100$, we can securely run 33 rounds before re-keying. One of the advantages of PCL is that it only relies on the DDH assumption in some cyclic group $\mathbb{G}$ of prime order p, like single-round KDK. However, there are also benefits regarding performance over multi-round KDK as highlighted in Section 4.

It is observed in [26] that KDK is centered on a fixed matrix A with entries in $\{-1, 0, 1\}$ that determine the exponents used to compute w. In other words, party P_i raises P_j's public key u_j to the power $A_{i,j}$ when computing w. However, the matrix A used in KDK has the form: $A_{i,i} = 0$ for $1 \leq i \leq 0$; $A_{i,j} = 1$ for $1 \leq i < j \leq n$ ($+1$ in upper triangle) and $A_{i,j} = -1$ for $1 \leq j < i \leq n$ (-1 in lower triangle). Therefore, A is skew-symmetric i.e. $-A = A^T$. The main idea in PCL is to generate a new skew-symmetric matrix $A^{(k)}$ in a deterministic manner for each round k. Furthermore, the matrix $A^{(k)}$ is chosen to have coefficients in $\mathbb{Z}_p$ instead of $\{-1, 0, 1\}$ in order to prove security. We refer the reader to [26] for details. Here we assume a function $\chi : \mathbb{Z}_p \times \mathbb{Z} \to \mathbb{Z}_p^{n \times n}$ that takes a random seed and a round number, and outputs a pseudorandom s kew-symmetric matrix over $\mathbb{Z}_p$. Note that the seed can be pre-determined or derived from the users' public keys. The main differences to single-round KDK are as follows:

- Let $s \in \mathbb{Z}_p$ be a seed deterministically derived form $u_1, \ldots, u_n$.
- In round k, party P_i computes
  - $A^{(k)} \leftarrow \chi(s, k) \in \mathbb{Z}_p^{n \times n}$.
  - $w \leftarrow \prod_{j \in 1}^{n} u_j^{A_{i,j}^{(k)}} \in \mathbb{Z}_p$.

The remaining steps are the same as single-round KDK with the exception that each party may choose a different input value in every round, and the final output is ℓ values in $\{0, \ldots, n\beta\}$.

4 Comparing the performance of multi-round KDK and PCL

We begin by comparing the *original* multi-round KDK from [20] and *original* PCL protocol form [26]. We are unaware of any concrete performance results for multi-round KDK, which we believe are important in order to assess its practicality. In this section, the performance of both multi-round protocols is measured and compared. The results motivate our proposed optimizations.

4.1 Computation of a round

Firstly, we compare the necessary group operations that a party P_i must perform in a given round. Multi-round KDK requires $n-1$ pairings, $n-1$ multiplications

in $\mathbb{G}_T$ and an exponentiation in $\mathbb{G}_T$. Note the omission of the inversions in $\mathbb{G}_T$ for $1 \leq j < i$. The reason for this is that in the *Setup* phase, party P_i can compute $U_j \leftarrow -U_j$ for $1 \leq j < i$ where $U_j \in \mathbb{G}_1$ is P_j's public key. Thus by bilinearity of e, no inversions are needed in $\mathbb{G}_T$.

On the other hand, PCL needs n exponentiations and n multiplications in group $\mathbb{G}$, Derivation of the *per-round information* for KDK involves computing $Q_k \leftarrow H_2(k) \in \mathbb{G}_2$ whereas PCL involves computing $A^{(k)} \leftarrow \chi(s, k)$. However, the latter can be optimized since only a single row of the matrix $A^{(k)}$ is needed by party P_i. As pointed out in [26], χ uses a hash function $H : \mathbb{Z}_p \times \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}_p$ to generate $A^{(k)}_{i,j}$; that is, $A^{(k)}_{i,j} \leftarrow H(s, k, i, j)$. However, derivation of the *per-round information* in both protocols is negligible relative to the cost of the group operations.

At present, all known efficient realizations of bilinear pairings are based on elliptic curves. Therefore, in order to implement multi-round KDK, we had to instantiate $\mathbb{G}_1$ and $\mathbb{G}_2$ with elliptic curve groups. There is far less freedom when choosing an elliptic curve when pairings are involved, since the chosen curve must satisfy additional properties. Notwithstanding, to provide a fair performance comparison between both protocols, the same curve was used for both protocols.

Consider an elliptic curve E over $\mathbb{F}_q$ for prime q whose order is $\#E(\mathbb{F}_q) = p$. For PCL, the group $\mathbb{G}$ may be instantiated by the additive group formed by $E(\mathbb{F}_q)$. For multi-round KDK, we have opted to use the Modified Tate Pairing to instantiate e since efficient implementations exist in libraries such as MIRACL.

Now the embedding degree k of E is the smallest positive integer such that $p \mid q^k - 1$. Concretely, the Tate pairing takes two points on $E(\mathbb{F}_{q^k})[p]$ and outputs an element of $\mathbb{F}^*_{q^k}$ (more precisely, an element of a multiplicative subgroup of order p of $\mathbb{F}^*_{q_k}$), where $E(\mathbb{F}_{q^k})[p]$ denotes the set of points on $E(\mathbb{F}_{q^k})$ of order p i.e. the set of points A with $pA = \mathcal{O}$, where $\mathcal{O}$ is the additive identity (*point at infinity*). Basically, $\mathbb{G}_1$ and $\mathbb{G}_2$ must be two distinct subgroups of $E(\mathbb{F}_{q^k})$ of order p. In fact, we can set $\mathbb{G}_1$ to $E(\mathbb{F}_q)$ to make the pairing calculation faster. Furthermore, certain pairing-friendly curves E allow us to choose $\mathbb{G}_2$ such that it is isomorphic to a subgroup of $E'(\mathbb{F}_{q^{k/d}})$ where E' is related to E (known as the "twisted curve"); this means arithmetic operations can be carried out in the smaller field $\mathbb{F}_{q^{k/d}}$ where d is the "twist degree".

The curve chosen for our implementation is a member of the pairing-friendly BN family from [27] with a 254-bit prime q, embedding degree $k = 12$, and "twist degree" $d = 6$. As a consequence of the latter, arithmetic operations in $\mathbb{G}_2$ can be carried out in $\mathbb{F}_{q^2}$ instead of $\mathbb{F}_{q^{12}}$. In addition, $\mathbb{G}_T$ is the group generated by $g = e(P, Q) \in \mathbb{F}^*_{q^k}$, and thus its arithmetic operations are carried out in the "big" field $\mathbb{F}_{q^k}$. This has implications for message recovery, since Pollard's lambda algorithm is much slower.

We implemented both protocols in C++ using the MIRACL C/C++ library version 5[1] using the BN curve as described above. The code was compiled with g++ with the compiler flags "-m64 -O2" as recommended in the MIRACL doc-

[1] On Github, commit `https://github.com/CertiVox/MIRACL/commit/6d7bb13285e7962ccfa110b4149fa8a63db2ed52`

umentation, and was executed on a machine with an Intel Core i5-3340M CPU (2.7 GHz) and 4GB of RAM, and running 64-bit Debian GNU/Linux 3.2.41. For each protocol. we measured the time taken to compute a single round per participant (recall that a round involves preparing the value v_i for some party P_i). We ran this 100 times for different values of n. Note that on each run a random index $i \in \{1, \ldots, n\}$ was chosen, and the round was executed for P_i. Our results are shown in Table 4.1. As expected, the cost of a round is roughly linear in n. Moreover, the difference in times between KDK and PCL is significant; on average PCL outperforms KDK by a factor of ≈ 437 based on the times in Table 4.1). For even a moderate number of users such as $n = 100$, it is clear that multi-round KDK is not suited to time-sensitive applications. This is more pronounced for resource-constrained devices such as wireless sensors.

	n=10	**n=100**	**n=1000**
Multi-round KDK	47.78 (0.25)	480.71 (2.34)	4795.33 (6.69)
PCL	0.94 (0.053)	1.33 (0.01)	5.33 (0.07)

Table 1. Mean time in ms (over 100 runs) for a party to compute a round (standard deviation in parenthesis).

4.2 Recovery of the sum with Pollard's Lambda Algorithm

Now we turn our attention to the aggregation phase of the protocol. In any given round, this entails multiplying all elements v_i to calculate $z \leftarrow \prod v_i$, then finding a discrete logarithm with respect to a generator g to recover the sum $\sigma = \sum m_i$ of the parties' inputs in the round. For this purpose, Pollard's Lambda algorithm is employed. In PCL, the v_i belong to $\mathbb{G}$ whereas in multi-round KDK, they belong to $\mathbb{G}_T$. Recall that our implementation with elliptic curves instantiates $\mathbb{G}$ as $E(\mathbb{F}_q)$ whereas $\mathbb{G}_T$ is instantiated as a subgroup of $\mathbb{F}^*_{q^k}$. As such, this phase is more expensive for multi-round KDK because the field operations take place in a "bigger" field. Pollard's Lambda algorithm dominates recovery of the sum. Its time complexity is $O(\sqrt{M})$ where M denotes the size of the message space. In this case, $M = n\beta$ since each party chooses her message in $\{0, \ldots, \beta\}$.

In order to compare multi-round KDK to PCL in this phase, we measured the time taken to compute Pollard's Lambda algorithm for different message space sizes. Moreover, values were randomly generated in the set $m \xleftarrow{\$} \{2^{b-2}, \ldots, 2^b\}$ for different values of b and the time taken to recover m given g^m using Pollards Lambda algorithm was measured (the range given to the algorithm was $\{0, \ldots, 2^b\}$); here g denotes the generator of the group in question and multiplicative notation is employed arbitrarily. The measurements were performed in Sage version 5.9 on the same machine and operating system as that used for the previous experiment above. We ran the experiment 10 times each for $b \in \{5, 10, 15, 20, 25, 30\}$ for both $E(\mathbb{F}_q)$ and the group $\langle e(P, Q) \rangle \subset \mathbb{F}^*_{q^k}$ (recall that $k = 12$ for the curve we used). Our results are shown in Figure 1. Observe that for ≈ 30-bit numbers, multi-round KDK takes almost half a minute to re-

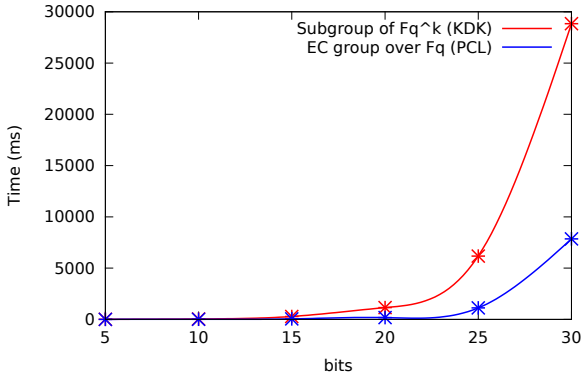

Fig. 1. Average time to find discrete logs in $E(\mathbb{F}_q)$ (PCL) and $\mathbb{F}_{q^k}$ (multi-round KDK) for different value ranges (upper bound in bits).

cover the result. Hence, for large values of β, it is the recovery phase that acts as the main bottlekneck in multi-round KDK.

5 Reducing Computational Load in the PCL Protocol

Although the performance results from Table 4.1 are favorable for PCL, there is still a motivation for seeking improvements, because many applications involve running the protocol on low-powered devices such as mobile phones. In particular, when a user's phone has low battery, it would be desirable to reduce their computational load. We introduce an optimization that allows parties to adjust their computational workloads at the expense of reducing their protection against collusion tolerance. Consider a skew-symmetric matrix A generated in a particular round of the protocol. Ordinarily, each row of the matrix has $n-1$ non-zero entries with overwhelming probability. It is easy to see that the cost of computing $v_i^{(k)}$ for some round k is linear in the number of non-zero entries in the associated matrix A. A user P_i can reduce this cost by setting some entries in row i of A to zero. Suppose he sets α entries of the i-th row of A to zero; the zero at $A_{i,i}$ is not included. His computation time is now $\frac{n-1-\alpha}{n-1}$ of the original. For convenience, we refer to these zeros as *holes*. If P_i sets a *hole* at position j, then P_j necessarily has a hole at position i. So if one party sets α holes, another α parties enjoy a marginally reduced cost. This idea raises some natural questions including

1. Since the matrix A is deterministically and non-interactively generated in each round, how are a party's holes set?
2. What impact does the number of holes α have on the collusion tolerance of the protocol? Does this lead to fewer rounds, and how many?

3. For some collusion tolerance t and some desired number of rounds m, how many users h can simultaneously set α holes while maintaining privacy?

The first question can be addressed by assuming that all parties who seek to place holes declare this intention in the initial stage of the protocol. We define the predicate $\mathfrak{h} : \{1, \ldots, n\} \to \{0, 1\}$ such that for all $1 \leq i \leq n$, we have $\mathfrak{h}(i) = 1$ if and only if P_i sets holes. For all parties P_i with $\mathfrak{h}(i) = 1$, we assume they all set α holes. The following algorithm is used to compute A in a given round. Note that we are using a hash function H with seed s effectively as a PRNG[2]. As such, we use the notation $x \xleftarrow{\mathsf{PRNG}(H,s)} X$ to denote the fact that x is pseudorandomly sampled from the set X using a pseduorandom function derived from H and s.

1. Set A to the $n \times n$ zero matrix.
2. Create an array W of length n; set $W[i] \leftarrow \alpha$ for every $1 \leq i \leq n$.
3. For all $1 \leq i \leq n$:
   (a) If $\mathfrak{h}(i) = 0$, set $S \leftarrow \{i + 1, \ldots, n\}$.
   (b) Else;
      i. Set $\alpha' \leftarrow W[i]$.
      ii. Choose a subset of $\{i + 1, \ldots, n\}$ of cardinality α' (these correspond to the holes). Formally, sample $\hat{S} \xleftarrow{\mathsf{PRNG}(H,s)} \{X \subseteq \{i + 1, \ldots, n\} : |X| = \alpha'\}$.
      iii. Set $W[j] \leftarrow W[j] - 1$ for every $j \in \hat{S}$.
      iv. Set $S \leftarrow \{i + 1, \ldots, n\} \setminus \hat{S}$.
   (c) For each $j \in S$:
      i. Set $A_{i,j} \xleftarrow{\mathsf{PRNG}(H,s)} \mathbb{Z}_p$.
      ii. set $A_{j,i} \leftarrow -A^{i,j}$

It turns out the all parties can set α holes and benefit from a performance boost. The following lemma gives a *lower bound* on the number of non-holes a party must set, and therefore the number of exponentiations in $\mathbb{G}$ (viewed multiplicatively) that have to performed per round (per party) to maintain t-privacy in the worst-case (this is the number of non-holes plus 1). Recall that n exponentiations are needed per round in the original PCL protocol.

Lemma 1. *Let ℓ be the number of rounds. Let $t < n$ be the collusion tolerance. A lower bound on the number of non-holes a party must set per round to guarantee t-privacy is $2\ell + t$.*

Proof. To derive this lower bound, we consider the worst-case scenario. This corresponds to the case where all parties choose the same α positions, which all lie outside the t parties controlled by the adversary.

A skew-symmetric matrix $A \in \mathbb{Z}^{n \times n}$ can be viewed as n quadratic equations in n variables $x_1, \ldots, x_n$. Moreover, the i-th row of A represents the equation

[2] A cryptographically secure pseudorandom number generator is not used directly because the proof of security is in the random oracle model, which relies on H being modelled as a random oracle.

$\sum_{j=1}^{n} A_{i,j}x_ix_j$. There are $n(n-1)/2$ unique monomials x_ix_j. There is a corresponding coefficient matrix $A' = \mathsf{coeff}(A) \in \mathbb{Z}^{n \times n(n-1)/2}$ for A whose i-th row consists of the coefficients for each monomial. By construction, A' is linearly dependant since all n equations must sum to zero to achieve correctness. But, if we remove one row (say the n-th) to yield A'', then A'' should be linearly independent. This is necessary to ensure security (see the proof of Theorem 1 in [26] for more details). If we have ℓ rounds, then the resulting coefficient matrix $C \in \mathbb{Z}_p^{\ell(n-1)\times(n-1)/2}$ (formed by vertically concatenating the first $n-1$ rows of $\mathsf{coeff}A^{(k)}$ for $k \leq \ell$) must be linearly independent to guarantee security. But by dropping the adversary's t parties, along with α parties corresponding to the holes (in the worst-case), we are left with $n-t-\alpha$ parties. A precondition for linear independence is that $\ell \leq \frac{n-t-\alpha}{2}$. Since each party must set $t+((n-t)-\alpha)+1$ non-holes, it follows from the inequality that a lower bound is $2\ell + t$. □

Lemma 1 only considers a lower bound. This tell us the best we can hope for. We explain in the next section why setting the number of non-holes to merely meet this lower bound is not sufficient for t-privacy.

5.1 Partitioning of the graph

Let A be a skew-symmetric matrix. A can be viewed as an undirected graph G, where the vertices represent the parties $P_1, \ldots, P_n$, and there is an edge between P_i and P_j if $A_{i,j} \neq 0$. If G can be partitioned into more than one connected component, say components G'_1 and G'_2, then partial sums can be learned of the parties in both components. It turns out that by just setting non-holes to merely meet the lower bound given by Lemma 1 is not sufficient to avoid partitions. Since A is deterministically generated, it might be tempting to modify the algorithm to generate A such that this does not occur. However, we don't know ahead of time which t parties are controlled by the adversary. Let m the number of non-holes set by each player. From Lemma 1, $m \geq 2\ell+t$. Formally, to ensure t-privacy, given a connected graph G of degree m, it must hold that if any $\kappa \leq t$ vertices are removed to yield G', then G' remains connected. In fact, such as graph G is referred to as k-vertex-connected, where k in this case is $t+1$. There are techniques to generate G to satisfy this property. One of those techniques involve each party linking to its m "nearest neighbors", where the distance between P_i and P_j is $|j-i|$. Therefore, we can change the algorithm that generates the skew-matrix A to follow this technique. This means that we can ensure t-privacy and perform only $2\ell + t + 1$ exponentiations per round, as opposed to n in the original protocol.

6 Reducing Computational Load in the KDK Protocol

Due to its comparitively poor performance, as shown in Table 4.1, there is abundant motivation for reducing the computational cost of a round of multi-round KDK. Our technique from Section 5 can also be applied to multi-round KDK,

although with even greater scope for improvement. The reason for this is as follows. Recall the coefficient matrix from the proof of Lemma 1. In PCL, $n-1$ rows are added to this matrix in every round. However, in multi-round KDK, due to the pairing, there is no linear relationship between the set of equations of each round. As such, we only have to consider a single set of $n-1$ equations in isolation. Since there are fewer equations, it is easier to avoid linear dependence, and more holes can be set as a consequence. Since a skew-symmetric matrix A is fixed for multi-round KDK, we replace this matrix with one generated according to the technique mentioned in Section 5.1.

Like above, we consider the worst-case scenario and derive a lower bound on α. Let us represent the number of dishonest users as a fraction $\tau \in [0,1]$ of n. As a necessary condition for linear independence, the following inequality must be satisfied

$$((1-\tau)n-\alpha)((1-\tau)n-\alpha) \geq 2((1-\tau)n-1). \tag{3}$$

Since we know that $\alpha < (1-\tau)n$, we can simplify (3) to obtain an upper bound on α:

$$\alpha \leq (\tau-1)n - 1/2\sqrt{8n-7} - 1/2. \tag{4}$$

Because there will always be τn non-holes, the number of additional non-holes that is necessary is $\lceil 1/2\sqrt{8n-7}+1/2 \rceil$. It follows that the computational load as a fraction of the original load is then lower bounded by

$$\frac{\tau \cdot n + \lceil 1/2\sqrt{8n-7}+1/2 \rceil}{n} = \tau + \frac{\lceil 1/2\sqrt{8n-7}+1/2 \rceil}{n} \tag{5}$$

which shows that as n grows, the cost of this modified protocol relative to the original converges towards τ.

7 Conclusions

While there are several protocols that enable users to privately compute the summary of their values, in many cases, as in the case of KDK, there are several hidden implementation bottlenecks which can significantly delay the calculations of different stages of the protocol. In this work we indicate that for instance the multi-round KDK protocol, due its heavy reliance on pairings is not practical at present for large numbers of users $n > 100$, whereas PCL is highly practical in these cases at the expense of reduced collusion tolerance. We showed that further customization of the privacy level facilitates further extensions to both protocols, and such extensions were shown to be secure. If applied, these extension can boost the efficiency of both these protocols, leading to faster applications with customizable levels of privacy.

References

1. Gilad Asharov and Yehuda Lindell. A full proof of the bgw protocol for perfectly-secure multiparty computation. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:36, 2011.

2. Paulo S. L. M. Barreto, Hae Yong Kim, Ben Lynn, and Michael Scott. Efficient algorithms for pairing-based cryptosystems. In Moti Yung, editor, *CRYPTO*, volume 2442 of *Lecture Notes in Computer Science*, pages 354–368. Springer, 2002.
3. Assaf Ben-David, Noam Nisan, and Benny Pinkas. Fairplaymp: a system for secure multi-party computation. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 257–266. ACM, 2008.
4. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 1–10. ACM, 1988.
5. Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas P Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, et al. Multiparty computation goes live. *IACR Cryptology ePrint Archive*, 2008:68, 2008.
6. Dan Boneh and Matthew Franklin. Identity-based encryption from the weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.
7. David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
8. David Chaum, Claude Crépeau, and Ivan Damgard. Multiparty unconditionally secure protocols. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 11–19. ACM, 1988.
9. Chris Clifton, Murat Kantarcioglu, Jaideep Vaidya, Xiaodong Lin, and Michael Y Zhu. Tools for privacy preserving distributed data mining. *ACM SIGKDD Explorations Newsletter*, 4(2):28–34, 2002.
10. I. Damgård, M. Geisler, M. Krøigaard, and J. Nielsen. Asynchronous Multiparty Computation: Theory and Implementation. *Public Key Cryptography*, pages 160–179, 2009.
11. I. Damgard, V. Pastro, N.P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. Cryptology ePrint Archive, Report 2011/535, 2011. `http://eprint.iacr.org/`.
12. Ratna Dutta, Rana Barua, and Palash Sarkar. Pairing-based cryptographic protocols : A survey. Cryptology ePrint Archive, Report 2004/064, 2004. `http://eprint.iacr.org/`.
13. Steven D. Galbraith, Keith Harrison, and David Soldera. In Claus Fieker and David R. Kohel, editors, *ANTS*, Lecture Notes in Computer Science, pages 324–337. Springer.
14. O. Goldreich. Towards a Theory of Software Protection and Simulation by Oblivious RAMs. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 182–194, 1987.
15. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 218–229. ACM, 1987.
16. Philippe Golle and Ari Juels. Dining cryptographers revisited. In *Advances in Cryptology-Eurocrypt 2004*, pages 456–473. Springer, 2004.
17. S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 513–524, New York, NY, USA, 2012. ACM.
18. Feng Hao and Piotr Zieliński. A 2-round anonymous veto protocol. In *Security Protocols*, pages 202–211. Springer, 2009.

19. Florian Hess, Nigel P. Smart, and Frederik Vercauteren. The eta pairing revisited. *IEEE Transactions on Information Theory*, 52(10):4595–4602, 2006.
20. Klaus Kursawe, George Danezis, and Markulf Kohlweiss. Privacy-friendly aggregation for the smart-grid. In *Privacy Enhancing Technologies*, pages 175–191. Springer, 2011.
21. Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. *IACR Cryptology ePrint Archive*, 2013:79, 2013.
22. Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *Advances in Cryptology-EUROCRYPT 2007*, pages 52–78. Springer, 2007.
23. Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. *Journal of cryptology*, 25(4):680–722, 2012.
24. Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. Fairplay-secure two-party computation system. In *USENIX Security Symposium*, pages 287–302, 2004.
25. Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In *Advances in Cryptology–CRYPTO 2012*, pages 681–700. Springer, 2012.
26. Constantinos Patsakis, Michael Clear, and Laird Paul. Private aggregation on untrusted servers with customizable thresholds. *IACR Cryptology ePrint Archive*, 2013:661, 2013.
27. Geovandro C. C. F. Pereira, Marcos A. Simplício, Jr., Michael Naehrig, and Paulo S. L. M. Barreto. A family of implementation-friendly bn elliptic curves. *J. Syst. Softw.*, 84(8):1319–1326, August 2011.
28. E. Shi, R. Chow, T. H. H. Chan, Dawn Song, and Eleanor Rieffel. Privacy-Preserving Aggregation of Time-Series Data. Technical report, UC Berkeley, 2011.
29. Z. Yang, S. Zhong, and R. N. Wright. Privacy-Preserving Classification of Customer Data Without Loss of Accuracy. In *SIAM International Conference on Data Mining*, pages 1–11, 2005.
30. Andrew Chi-Chih Yao. Protocols for secure computations. In *FOCS*, volume 82, pages 160–164, 1982.
31. Andrew Chi-Chih Yao. How to generate and exchange secrets. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*, pages 162–167. IEEE, 1986.